

*Report of the
Defense Science Board*

on

**TASK FORCE ON
DOD SUPERCOMPUTING NEEDS**



11 October 2000

*Office of the Under Secretary of Defense
For Acquisition and Technology
Washington, D.C. 20301-3140*



OFFICE OF THE SECRETARY OF DEFENSE

3140 DEFENSE PENTAGON
WASHINGTON, DC 20301-3140

DEFENSE SCIENCE
BOARD

MEMORANDUM FOR UNDER SECRETARY OF DEFENSE (ACQUISITION,
TECHNOLOGY AND LOGISTICS)

SUBJECT: Final Report of the Defense Science Board Task
Force on DoD Super Computing Needs

I am forwarding the final report of the Defense Science
Board Task Force on DoD Super Computing Needs.

The Terms of Reference directed the Task Force to
address DoD Super Computing Needs in light of recent
commercial marketplace developments. Specifically, the Task
Force was tasked to assess whether DoD should continue its
investment in the development of the CRAY SV2.

The Task Force formulated three recommendations which
address DoD near term, medium term, and far term needs while
taking into account the dynamic nature of the High
Performance Computing marketplace. I believe these
recommendations best position DoD to take advantage of the
benefits offered by the High Performance Computing industry
while mitigating its overall risk.

I endorse all of the Task Force's recommendations and
propose you review the Task Force Chairman's letter and
report.

A handwritten signature in black ink, appearing to read "Craig Fields", is positioned above the printed name.

Craig Fields
Chairman



OFFICE OF THE SECRETARY OF DEFENSE

3140 DEFENSE PENTAGON
WASHINGTON, DC 20301-3140

DEFENSE SCIENCE
BOARD

MEMORANDUM FOR CHAIRMAN, DEFENSE SCIENCE BOARD

SUBJECT: Final Report of the Defense Science Board Task Force on
DoD Super Computing Needs

Attached is the report of the Defense Science Board Task Force on DoD Super Computing Needs.

The Task Force was created as a spin off of a larger effort investigating Defense Software issues and was tasked to review DoD Super Computing Needs. Specifically, the Task Force was charged with examining DoD needs related to the field of cryptanalysis in light of emerging trends in the High Performance Computing market.

The Task Force validated the need for high performance computers that provide extremely rapid access to extremely large global memories. This capability would support not only cryptanalysis but several other important DoD needs as well (e.g. calculation of weapons effects, weapon design and analysis, acoustic analysis, computational fluid dynamics, radar cross sectional modeling, and synthetic materials design).

The Task Force recommends a three part strategy to meet the DoD's Super Computing Needs. First, the DoD should continue short-term support of the CRAY SV2 development. This is a risky development, but the modest expenditures are worth the potential payoff in performance improvement. Secondly, the DoD should develop a high bandwidth memory system using Commercial-off-the-Shelf microprocessors for the medium term. This strategy mitigates any potential failure of the SV2 development. Finally, DoD should invest in long-term research to address unique Defense computing needs. Such research is essential to refill the Research and Development pipeline with new technologies that will enable tomorrow's high performance computers.

The Task Force would like to express its appreciation for the cooperation, advice, and help by the government advisors, support staff, and the many presenters from commercial computing firms and research organizations.

Bob Nesbit

Mr. Bob Nesbit
Task Force Chairman

TABLE OF CONTENTS

| | |
|--|-----|
| Table of Contents | i |
| Executive Summary | 1 |
| Findings | 1 |
| Recommendations | 2 |
| Introduction | 4 |
| Background | 4 |
| Assessing the National Security Need | 5 |
| Assessing the Commercial HPC Market | 8 |
| Recommendations | 111 |
| Annex A. Briefings Received | A-1 |
| Annex B. Tasking Memorandum | B-1 |

EXECUTIVE SUMMARY

The Defense Science Board Task Force on Defense Software was asked to form a subgroup to examine changes in supercomputing technology and investigate alternative supercomputing technologies in the areas of distributed networks and multi-processor machines. The work of the Task Force was motivated by recent DoD investment decisions involving the development of next-generation High Performance Computers (HPC) to be used for cryptanalysis. The Task Force did not consider alternative investment strategies into other techniques besides code breaking.

Toward this end, the Task Force studied the DoD's need for HPC, assessed the HPC market as it affects the DoD and made recommendations for near, mid and long-term strategies that should be implemented in order to insure DoD's future HPC needs are met.

Findings

The Task Force concluded that there is a significant need for high performance computers that provide extremely fast access to extremely large global memories. Such computers support a crucial national cryptanalysis capability. To be of most use to the affected research community, these supercomputers also must be easy to program. It is also clear that the current mainstream commercial HPC market is not producing systems that meet this critical DoD need.

The Task Force determined that beyond cryptanalysis, the national security need for HPCs with high-global-memory bandwidth is not as widespread as it once was. Nonetheless, there are other national security applications that would likely benefit from the existence of a system providing high-global-memory bandwidth, including:

- calculation of weapons effects
- weapon design and analysis
- acoustic analysis
- computational fluid dynamics
- radar cross section modeling
- synthetic materials design

Our limited study did not have a chance to assess and validate in depth any threat to national security of not being able to support these applications in the future.

An important consideration in the Task Force's deliberations was the assessment of the overall HPC market, market directions, and the market potential for supporting the continued development of traditional high-global-memory-bandwidth vector supercomputers like the Cray SV2 in the future.

The vector supercomputing portion of the capability segment of the high performance technical computing market is at a critical juncture as far as US national security interests are concerned. If the current Cray SV2 development slips its schedule or is unsuccessful, this vector market will be lost to the US with the result that only foreign (Japanese) sources will be available for obtaining this critical computing capability.

Vector supercomputing will continue to be pressured at the high-end by the large-scale parallel systems, and where vector machines hold sway, Cray will face stiff foreign competition in non-US markets. Unless the market situation changes significantly, there appears to be insufficient commercial demand for vector supercomputers to support the current number of vendors.

Recommendations

To meet the DoD need for supercomputers with high-global-memory bandwidth, the Task Force recommends that the DoD pursue a three-part strategy to ensure the supply and continued evolution of High Performance Computers. The three parts of the strategy are aimed at ensuring capability in the short term (within 2 years), the medium term (2 to 5 years), and the long term (beyond 5 years).

1. Support the development of Cray SV2 in the short term.

To meet DoD needs in the short term, the Task Force recommends that the DoD continue to support the development of the Cray SV2. This machine potentially will be capable of two orders of magnitude more global-memory bandwidth than today's T-90 or T3E as well as tomorrow's cluster-based machines available from commercially mainstream HPC vendors. We see little possibility of any other vendor being able to deliver a machine with this capability within the next two years.

While the Task Force considers the development of the SV2 to be a *very* high-risk venture, we believe the DoD should continue to pursue its development because the potential payoff is so great – two orders of magnitude improvement – and the required investment is reasonable.

It should be understood that supporting the SV2 might not be a one-time expense but rather a continuing investment in a critical defense-specific capability. At present, there appears to be insufficient commercial demand for this class of machines to make this industry self-supporting. Unless the market situation changes significantly, continued investment will be necessary to support the further evolution of vector supercomputers.

2. For the medium term, develop an integrated system based on COTS microprocessors and a new high-bandwidth memory system.

Because of concerns associated with the ongoing development of the SV2, the Task Force recommends this second option be initiated and pursued in parallel to reduce the national security risk of being without a future organic high-global-memory-bandwidth computing capability. The bandwidth needs of critical DoD applications can be met without the expense or loss of scalar performance associated with building a custom vector processor. COTS microprocessors can be leveraged for these applications by building a very-high-bandwidth memory system. We expect it is feasible to build such an integrated system with a global-memory bandwidth three orders of magnitude higher than the T3E. However, there are significant risks associated with the difficulty of programming such an integrated system that need to be addressed along the way to assure its ultimate usefulness to the research community.

Depending on the degree of success on the targeted cryptanalysis application of the SV2 or the microprocessor-based integrated system, the DoD will have the option in the future to continue evolving the SV2 line or switching to and maturing the integrated system. This later

case will almost certainly require continued DoD investment in the future as we believe it is unlikely that the integrated system will be commercially viable on its own.

The National Security Agency (NSA) and Director Defense Research and Engineering (DDR&E) are jointly sponsoring the development of the SV2. Funding and direction for development of this alternative integrated system using COTS microprocessors could be similarly a joint effort. But to simplify the situation, we suggest that it is more reasonable for NSA to focus on the SV2 and DDR&E to undertake the COTS microprocessor-based integrated system.

3. Invest in research on critical technologies for the long term.

The third recommendation of the Task Force is for the DoD to invest in long-term research to address unique Defense computing needs. For the performance of high-global-memory-bandwidth systems to continue to scale, long-term research is essential to refill the Research and Development (R&D) pipeline with new technologies that will enable tomorrow's supercomputers.

Research investments should be made in strategic technologies that are critical to high-performance computing but are not being addressed by commercial industry. Important research areas include:

- architecture of high-performance computer systems
- memory systems, and I/O systems
- high-bandwidth interconnection technology
- system software for high-performance computers
- application software and programming methods for high-performance computers.

Research of this type, as opposed to development, is best carried out by universities and research laboratories where scientists can focus on long-term research without the pressing need to support short-term development.

INTRODUCTION

The Defense Science Board (DSB) was asked to examine changes in supercomputing technology and investigate new supercomputing alternatives for the Department of Defense – especially as related to the field of cryptanalysis¹. The terms of reference dated 15 November 1999 is provided in Annex B.

A DSB Task Force on High Performance Computing was formed with the following members: Dr. William J. Dally, Stanford University; Dr. Richard Games, MITRE; Mr. Robert Graybill, DARPA; Dr. Robert F. Lucas, Lawrence Berkeley National Laboratory; and Mr. Robert Nesbit, MITRE, who served as chairman of the group. Dr. Charlie Holland was the OSD point of contact. LtCol David Luginbuhl, USAF, served as executive secretary and CDR Brian Hughes, USN, the DSB secretariat representative. Dr. William Carlson from the Institute for Defense Analysis attended several meetings and provided valuable insights on certain technical matters.

The Task Force held four two-day meetings. The first in December 1999 at the National Security Agency to discuss their specific HPC needs, programs and plans. Also at that meeting SGI/Cray presented the SV2 design and progress. The second meeting in February 2000 was held in Washington to review numerous other DoD, government, and commercial HPC applications. In the third session in March 2000 at Lawrence Berkeley National Laboratory we met with six HPC vendors – Sun, HP, Mercury, IBM, Fujitsu, and Compaq – to discuss their future product plans. The final meeting in May 2000 included a presentation on HPC market trends as viewed by the International Data Corporation, an update on the DoE Accelerated Strategic Computing Initiative, and a discussion of the “new” Cray Inc. with their CEO James Rottsohlk. Tera Computer purchased the Cray division from SGI during the course of the study and adopted the Cray name. Annex A provides more details on the briefings the Task Force received.

The work of the Task Force was motivated by recent DoD investment decisions involving the development of next-generation supercomputers to be used for cryptanalysis. The Task Force did not consider alternative investment strategies into other techniques besides code breaking.

Our observations, findings and recommendations were discussed with Director, Defense Research and Engineering, Dr. Hans Mark, and Deputy Under Secretary of Defense (Science and Technology), Dr. Delores Etter on 5 May 2000. This letter summarizes and documents the work.

BACKGROUND

The market for the highest performance computing systems is relatively small. The National Security community within the US government has always been the largest customer for high performance computers, especially the high-global-memory-bandwidth systems available in the past from companies like Cray Research. During the last decade, pressures on US Defense budgets have significantly reduced the market for these very high performance systems. While

¹ Although the terms of reference specified “cryptography” (making of codes) it became apparent that it was the cryptanalysis application that was the real motivation for the study.

there has been some growth in the commercial market for such systems, it is not enough for the overall market to grow.

At the same time as the Defense market began shrinking, a number of competitors tried to enter the high performance computing market. These included Japanese companies with vector mainframes as well as a new generation of US companies offering scalable systems based on commodity microprocessors. This was driven in part by technology and in part by government investment. The Ministry of International Trade and Industry (MITI) pushed vector investments in Japan. The Defense Advanced Research Projects Agency (DARPA) put its investment money into scalable computing. More recently, the Department of Energy (DOE) ASCI program has led US R&D investments in scalable machines. The net result was the fragmentation of the high-end marketplace into an environment where no companies were profitable. Large vertical companies such as NEC and Fujitsu absorbed the losses. Smaller companies such as Thinking Machines, Kendall Square and Encore went bankrupt. And while Cray Research was acquired by Silicon Graphics, Inc. (SGI), there was little investment made by the company in new vector supercomputer developments.

The high performance computing marketplace has further been squeezed by the increasing performance of smaller workstations and servers. Large supercomputers have always been the only way to solve some really big, "capability" problems. In the past they were also the most cost-effective way to provide the "capacity" to address a multitude of smaller problems. Much of this "capacity" workload has moved in the last decade to workstations, servers, and even PCs, which have become the most cost-effective platforms. We discuss these market trends in more detail later in the report.

Recent scalable systems consist of networked compute nodes, each with their individual memory, and have sacrificed memory bandwidth in the quest for maximum cost-effectiveness. The result is that scalable systems have performance problems with global scatter/gather and irregular memory access patterns that vector machines traditionally have performed well on. Also the distributed-memory model of scalable systems is more difficult to program than the shared-memory model of past vector machines. Past vector machines from Cray Research have been relatively easy to use, and this has allowed the research community to get preliminary results quickly and without the need to optimize algorithms or code.

ASSESSING THE NATIONAL SECURITY NEED

The Task Force concluded that there is a significant national security need for high performance computers that provide extremely fast random access to a large global memory. It was also clear that the current mainstream commercial HPC market is not producing systems that meet this need. In the past supercomputers produced by Cray Research have featured the desired high-global-memory bandwidth, as well as specialized vector processors useful in some applications. However, mainstream commercial HPC systems today incorporate commodity microprocessors coupled to cheaper and less capable memory subsystems that provide significantly slower global-memory access rates.

The Task Force determined that the cryptanalysis application domain has a critical requirement for HPCs with high-random-access-global-memory bandwidth. There are three dimensions to this computing requirement:

- (1) the rate of random access to global memory measured in billions of updates/second (GUPS)
- (2) the size of the global memory, and
- (3) the ease of programming.

The first two dimensions translate directly into application capability. The third dimension bears on how easy it is to actually apply the computing capability. In the case of research activities involving a domain expert, even one with significant computer science skills, a difficult programming environment can eliminate an otherwise capable system from consideration. Ease of programming is also important for operational uses, but it usually does not represent a "show stopper" since application programs can be built to specification by a team of expert programmers. Table 1 summarizes the current situation along these three requirement dimensions for various classes of current and proposed HPC architectures. Actual benchmarked GUPS values for 4 GB tables are also shown.

Table 1. Three Dimensions of Computing Capability

Key: green = . provides the most useful capability (today)
 yellow = provides a marginal capability (today)
 red = provides only a limited capability (today)

| Architecture (Year) | GUPS (4GB) | Memory Size | Programmability |
|-------------------------------------|------------------------|--------------------|------------------------|
| <u>Parallel Vector</u> | | | |
| Cray YMP (1988) | red (.16) | red | green |
| Cray C90 (1991) | yellow (.96) | red | green |
| Cray T90 (1995) | yellow (3.2) | red | green |
| Cray SV1 (1999) | yellow (.7) | yellow | green |
| <u>Massively Parallel Processor</u> | | | |
| Cray T3E (1996) | yellow (2.2) | green | yellow |
| <u>Symmetric Multiprocessor</u> | | | |
| Multiple Vendors | red/yellow (.35 - 1) | yellow | green |
| <u>Clusters</u> | | | |
| Multiple Vendors | red/yellow (.35 - 1) | green | red |
| <u>Scalable Vector</u> | | | |
| Cray SV2 (2002) | green (400 govt. est.) | green | yellow |

Table 1 demonstrates that there has not really been a significant improvement in the GUPS measure of global-memory bandwidth since the factor of six increase at the transition from the Cray YMP to the Cray C90, which occurred in 1992. In fact the recent trend is that mainstream commercial symmetric multiprocessors (SMPs) and clusters are providing less GUPS capability. The scalable MPP and cluster systems do provide massive amounts of memory, but they are more difficult to program. An example of this is the Cray T3E, which has a well-engineered memory system that provides a GUPS rating on par with the Cray T90, but because of its different programming model has had less research impact in the application domain. The proposed Cray SV2 system is expected to provide a GUPS rate that is orders of magnitude higher than any system available today as well as a total memory size on par with scalable cluster systems. However, programming the SV2 will be more difficult than previous parallel vector systems because of its non-uniform memory access rates.

What about the non-commercially supported HPC national security needs beyond that of cryptanalysis? The national security need today for HPCs with high-global-memory bandwidth is not as widespread as it once was. This is because a large number of national security applications have been retooled or have been developed from the start to run on high-end commercial servers or clusters. Most notable in this retooling effort is the DOE Accelerated Strategic Computing Initiative (ASCI) program for nuclear stockpile stewardship and a variety of efforts supported by the DoD HPC Modernization program. The performance of these retooled codes depends on the application's communication requirements – a lot of fine-grain, random, global-memory accesses will especially degrade performance. This retooling has narrowed the size of the future national security market for high-global-memory-bandwidth HPCs.

Nonetheless, there are other national security applications that would likely benefit from the existence of a system providing high-global-memory bandwidth. Many of these are scientific and engineering applications that require implicit solutions of partial differential equations discretized on irregular grids. Examples include calculation of weapons effects, the design and analysis of weapons and platforms, acoustic analysis of submarines and computational fluid dynamics. Other applications include radar cross section modeling and designing synthetic materials. Our limited study did not have a chance to assess and validate in depth any threat to national security of not being able to support these applications in the future.

The Task Force also heard about commercial and civilian research applications (e.g. structural analysis, crash codes, climate modeling, and quantum chemistry) that benefit from the high performance delivered by the vector processors of a traditional high-global-memory-bandwidth supercomputer. Some presenters suggested implications to the United States' industrial competitiveness if access to future vector supercomputers was not assured, but this topic was beyond the scope of our Task Force.

In summary, there is a significant, albeit somewhat narrow, need for high performance computers that provide extremely fast access to extremely large global memories. Such computers support a crucial national cryptanalysis capability. To be of most use to the affected research community, these supercomputers also must be easy to program.

ASSESSING THE COMMERCIAL HPC MARKET

An important consideration in the Task Force's deliberations was the assessment of the overall HPC market, the market directions, and the market potential for supporting the continued development of traditional high-global-memory-bandwidth vector supercomputers like the Cray SV2 in the future. Using the IDC market definitions, the overall high performance technical computing market may be divided into four segments: 1) Technical Capability, 2) Technical Enterprise, 3) Technical Divisional, and 4) Technical Departmental. The first market segment, traditionally viewed as the high-end supercomputing or HPC market, is driven by a relatively small number of users with large specialized applications requiring high-end computing capability. Typically a single program may consume an entire computing system.

The other three technical computing markets segments are driven to a larger degree by a large number of end users with lots of small jobs that run simultaneously on a multiple-user machine or on many single-user machines. As such, these three market segments can be grouped together and referred to as the *technical capacity* market, where the throughput delivered on many small jobs is the important metric. The technical capacity market is dominated by commodity microprocessor-based systems from Compaq, HP, IBM, SGI, and Sun. These same systems, mostly various-sized SMP systems, are also sold into the much higher volume commercial database market, providing these companies with a broad base to support continued research and development of next generation systems.

The total worldwide high performance technical computing revenue for 1999 was estimated by IDC to be \$5,617M. This breaks down to \$934M for the high-end technical capability market and \$4,683M for the technical capacity market. Figure 1 shows the worldwide trends in total revenues according to IDC for the high-end technical capability and technical capacity markets over the last five years. The technical capacity market has grown significantly while the high-end technical capability market has been fixed at around \$1,000M. Some traditional high-end users are moving down a segment because of increased computational capability offered at lower segments.




Figure 1. Technical Capacity versus Technical Capacity Revenue Comparison

Over the last 10 years, the technical capability market has expanded beyond just the traditional vector supercomputers to include large-scale parallel computing platforms based on commodity microprocessors. These platforms include the massively parallel processors (e.g., Cray T3E or Intel Paragon/ASCI Red) or large networked clusters of commercially mainstream SMPs from multiple vendors. We noted previously the DoE and DoD software retooling efforts that have helped to shift market share away from the vector supercomputers to large-scale parallel systems. According to IDC the total high-end technical capability revenue of \$943M for 1999 is divided into sales of \$500M for traditional vector supercomputers and \$443M for large-scale parallel HPCs.

Figure 2 focuses only on the vector supercomputing segment of the high-end technical capability market and shows the worldwide revenue trends according to IDC for the last five years. This market in total has remained relatively constant at about \$500M over this period. But there has been a dramatic shift in market share with the Japanese vendors currently dominating this market segment. The most significant factor that contributed to the decline in US market share in this segment is that Cray, while a division of SGI, did not produce a vector supercomputing product generation that can compete effectively with current Japanese offerings. A second factor is the aggressive pricing by the Japanese vendors. This can be addressed in the US by trade policy but poses a future challenge for Cray as it attempts to regain market share in Europe with its forthcoming SV2 system. Market share in the long term enables a company to generate the large returns required to develop the next generation of high-end computers and remain competitive in this critical but rather high development cost business.




Figure 2. Vector Supercomputer Revenue

What are the market projections for the future? IDC projects that by 2003 the technical capacity market will grow from the current \$4,683M to \$6,300M (compound annual growth rate of 9.3%), while the technical capability market will grow from the current \$934M to \$1,200M (6.7% compound annual growth rate (CAGR)). It remains to be seen to what extent the class of vector supercomputers, and the Cray SV2 in particular, will participate in this projected modest

market growth of the technical capability segment remains. One possible source of additional demand is the increasing emphasis on computer-aided engineering in the automotive and aerospace markets. Additionally, there is a possibility of emerging markets for traditional vector supercomputers in biotechnology and database processing (e.g., credit card fraud detection) applications.

In summary, the vector supercomputing portion of the capability segment of the high performance technical computing market is at a critical juncture as far as US national security interests are concerned. If the current Cray SV2 development slips its schedule or is unsuccessful, this vector market will be lost to the US with the result that only foreign (Japanese) sources will be available for obtaining this critical computing capability. Even if Cray can execute the development of the SV2 as planned, the road ahead will still be a difficult one. Vector supercomputing will continue to be pressured at the high-end by the large-scale parallel systems, and where vector machines hold sway, Cray will face stiff foreign competition in non-US markets. Unless the market situation changes significantly, there appears to be insufficient commercial demand for vector supercomputers to support the current number of vendors. Further discussion on this topic and how to respond is included in the Task Force's recommendations.

RECOMMENDATIONS

To meet the need for supercomputers with high-global-memory bandwidth we recommend that the DoD pursue a three-part strategy to ensure the supply and continued evolution of these machines. The three parts of the strategy are aimed at ensuring capability in the short term (within 2 years), the medium term (2 to 5 years), and the long term (beyond 5 years).

To place the suggestions that follow into context, we note that other US government agencies are aware of the limitations of today's commercial systems and are making modest investments to address these problems. The DoE ASCI Path Forward program is spending \$25M per year with IBM, Compaq, Sun, and others to address interconnect bandwidth and other deficiencies in SMP clusters. NASA is spending \$17M per year to get bigger SMP systems from SGI.

1. Support the Cray SV2 in the short term. To meet the need in the short term, we recommend that the DoD continue to support the development of the Cray SV2. This machine potentially will be capable of two orders of magnitude more global-memory bandwidth (GUPS) than today's T-90 or T3E as well as tomorrow's cluster-based machines available from commercially mainstream HPC vendors. We see little possibility of any other vendor being able to deliver a machine with this capability within the next two years.

The DoD should ensure that the Cray SV2 is completed by the end of 2002 by continuing to directly fund a portion of the development, by being a good customer, and by closely monitoring the project. By being a *good customer*, that is providing letters of intent or purchase orders for a regular stream of machines, the DoD can enhance Cray's ability to raise the capital needed to fund the project on the private equity markets. By closely monitoring the project, the DoD can increase the probability of timely delivery, particularly in light of the concerns expressed below.

We have two concerns relating to the development of the Cray SV2: lack of focus, and poor performance on scalar code. Cray Inc., a small company with limited resources, is currently dividing its effort between two unrelated supercomputer development projects: the Cray SV2, and the Tera Multithreaded Architecture (MTA). Their probability of success, and in particular the probability of timely delivery, would be greatly enhanced if they could be persuaded to focus their efforts entirely on the SV2. For example, schedule risk could be substantially reduced if software resources currently assigned to the MTA could be redirected to the SV2 and if the size of the SV2 prototype build could be increased. A company the size of Cray needs to focus all of its efforts on a single architecture and a single supercomputer.

The scalar processor in the Cray SV2 is a relatively simple processor operating at a modest clock rate. We expect such a processor to have significantly lower scalar performance than a high-end commercial microprocessor such as a Compaq Alpha, IBM Power4, or Intel Itanium that have four to six-issue out-of-order pipelines that operate at clock rates of exceeding 1 GHz. While this lag in scalar performance does not directly impact DoD applications that depend on vector performance rather than scalar performance, it will make this machine much less attractive to many commercial users that run code that cannot be completely vectorized.

It should be understood that supporting the SV2 may not be a one-time expense but rather a continuing investment in a critical defense-specific capability. At present, there appears to be insufficient commercial demand for this class of machines to make this industry self-supporting. Unless the market situation changes significantly, continued investment will be necessary to support the further evolution of vector supercomputers.

Given all the technical, market, and organizational issues, we consider the SV2 development to be a *very* high-risk venture. The DoD should continue to pursue the development because the potential payoff is so great – two orders of magnitude improvement – and the required investment is reasonable. But considering the very high risk, it is extremely important to pursue an alternative approach. Our suggestion follows.

2. For the medium term, develop an integrated system based on COTS microprocessors and a new high-bandwidth memory system. Because of our concerns associated with the ongoing development of the SV2, the Task Force recommends this second option be initiated and pursued in parallel to reduce the national security risk of being without a future organic high-global-memory-bandwidth computing capability.

The bandwidth needs of critical DoD applications can be met without the expense or loss of scalar performance associated with building a custom vector processor. COTS microprocessors can be leveraged for these applications by building a very-high-bandwidth memory system. Such a system would employ COTS DRAM chips, ASIC memory controllers, a high-bandwidth interconnection network, and a latency-hiding processor interface similar to the E-registers on the T3E. We expect it is feasible to build such an integrated system with a global-memory bandwidth in excess of 1000 GUPS by 2003 – three orders of magnitude higher than the GUPS for the T3E.

This approach should be less expensive than developing a complete vector computer system since the cost of developing the vector processor, scalar processor, cache subsystem, and the software to support the processors is eliminated. Commercial microprocessors along with their operating systems and compilers may be used with a few modifications. For example, operating system and compiler extensions would be needed to support the very-high-bandwidth memory system. Moreover, this approach results in better scalar performance than a vector processor because it leverages the considerable commercial investment in high-performance microprocessor design. The DoD should also try to introduce compatible changes to future COTS processor designs (e.g., special instructions or concepts like processor in memory) to make the high-bandwidth memory system more effective.

A program to develop a high-bandwidth memory system of the type described here would be best undertaken by a company with expertise in interconnection networks, system integration with COTS processors, and in delivering reliable hardware systems. Examples of such companies include Quadrex and Mercury.

Furthermore, it is important that such a future integrated system be easy to program and come with state-of-the-practice software tools (e.g., compilers, debuggers, languages such as IDA's UPC, and the Message Passing Interface). Although certain COTS software components can be leveraged, providing a robust and usable system software environment for the integrated system is a non-trivial task and would take some further effort and time to mature. As a future goal this integrated system should be easier to program than today's counterpart—the T3E. In concert with pursuing this hardware strategy, software technologies that propose to make such a

future integrated system more accessible to researchers, such as IDA's UPC, should be demonstrated today. The T3E provides a test bed today for software technology improvements that can effectively engage current researchers. Therefore, the future use of UPC on the T3E should be encouraged and the results closely followed.

There is some risk that a highly capable integrated system of the sort described here would further fragment the high-end technical capability market, further pressuring vector supercomputers like the SV2 and any follow-on systems. The impact such an integrated system would actually have would depend on its commercial prospects beyond the intended national security applications. Because of the cost of the high-bandwidth memory system, it will be significantly more expensive than large-scale parallel clusters, but may compete with them on applications that are bandwidth limited.

This potential "market confusion" factor caused by the development of the integrated system needs to be explicitly managed as part of future DoD investment decisions. It is difficult to predict the future or address all the possibilities, but the following three major cases can be identified conditioned on the degree of Cray's success with the SV2:

Best Case: The SV2 development is successful and the wide applicability of vector processing results in market growth for this type of machine *and* Cray is able to capture a substantial share of this increased market size to support future developments. Then the need for continued government investment in Cray product development would decrease. This would also reduce the need of ongoing government investment to mature/evolve the integrated system.

Middle Case: The SV2 development is successful, but there is not sufficient growth in Cray's market share to sustain future Cray development without continuing government investment. Then the future government investment decision should also factor in the success of the integrated solution. If both options are successful, then one key discriminator for follow-on investment will be which one has engaged more effectively the targeted cryptanalysis research/application community.

Worse Case: The SV2 development falters. Then future near-term incremental DoD investments in Cray should be stopped, and the majority of the resources should be focused on making the integrated system a success. We don't think it is likely that the integrated system will be commercially viable, and so its evolution will most likely require continued DoD investment.

Pursuing both the SV2 and the integrated-system developments in parallel for the next two years will provide the DoD with the most options. We don't expect the best case scenario to occur, and so the integrated system becomes either a useful point of comparison (for the middle case) or *crucial* (for the worse case) depending on the future.

The NSA and DDR&E are jointly sponsoring the development of the SV2. Funding and direction for development of the alternative integrated system using COTS microprocessors could be similarly a joint effort. But to simplify the situation, we suggest that it is more reasonable for NSA to focus on the SV2 and DDR&E to undertake the COTS microprocessor-based integrated system.

3. Invest in research on critical technologies for the long term. The third recommendation of the Task Force is for the DoD to invest in long-term research to address unique Defense computing needs. There has been little long-term research on high-performance computing in recent years and the reservoir of high-performance computing techniques that has for years been

trickling down from mainframes and supercomputers to microprocessors is nearly at an end. For the performance of high-global-memory-bandwidth systems to continue to scale, long-term research is essential to refill the R&D pipeline with new technologies that will enable tomorrow's supercomputers.

Research investments should be made in strategic technologies that are critical to high-performance computing but are not being addressed by commercial industry. Important research areas include: architecture of high-performance computer systems, memory systems, and I/O systems; high-bandwidth interconnection technology (architecture, signaling technology, and packaging technology); system software (compilers, operating systems, I/O software, and programming environments) for high-performance computers; application software and programming methods for high-performance computers. Areas such as single-processor architecture and semiconductor technology that are adequately addressed by industry should not be the focus of such a program.

Research of this type, as opposed to development, is best carried out by universities and research laboratories where scientists can focus on long-term research without the pressing need to support short-term development. The program should focus research funding on a few areas with funding in each area sufficient to engage the top scientists and achieve a critical mass rather than spread funding thinly over many areas. Research should focus on technologies at an advanced stage where success is not yet assured. To mitigate risk, several high-risk approaches to each key problem should be pursued on a pilot scale with a plan to down select before proceeding to development.

ANNEX A. BRIEFINGS RECEIVED

**DEFENSE SCIENCE BOARD TASK FORCE ON
DoD SUPERCOMPUTING NEEDS
BIEFINGS RECEIVED**

13-14 December

Overview of NSA's High Performance Computing Program
NSA HPC – Operational Needs and Strategies
NSA General Crytanalytic Computing Topics and Case Histories
NSA High End Computers – Comparisons and Contrasts
CRAY Initiatives
NSA HPC Alternatives R&D Program
NSA Long Term R&D Programs

Mr. Steve Oberlin

3 - 4 February 2000

Lockheed Martin
National Institute of Health
Computational Requirements for Finite Element Analysis - Boeing
NASA Goddard

Mr. Gary Mastin
Dr. Stan Burt
Mr. Roger Grimes

Mr. Richard Rood

Defense Threat Reduction Agency

Mr. Gene Stokes
Ms. Jacqueline Bell
Dr. Dan Hitchcock

Department of Energy
DoD Common HPC Software Support Initiative:
Overview

DNS and LES of Turbulent Flows in Complex Geometry
Massively Parallel Simulations of Fluid-Structure Interaction
High Performance Computing Requirements for Computational Terminal Ballistics
Computational Assisted Development of High-Temperature Structural Materials
Computational Chemistry and Materials Science (CCM)

Mr. Cray Henry
Mr. John Grogh
Dr. George Karniadakis
Dr. Joseph Baum
Mr. Eric Mestreau
Mr. Stephen Schraml

Dr. Rajiv K. Kalia
Dr. Priya Vashishta
Dr. Jerry Boatz

30 - 31 March 2000

SUN

Mr. Robert Bredehoff
Mr. Nicholas Aneshansley
Mr. Steven Sistare
Mr. Timothy Morgan

Hewlett Packard
Mercury
Fujitsu

COMPAQ
IBM

4 - 5 May

IDC - HPC Market Trends

Cray Inc.

Task Force meeting with DDRE - Dr. Hans Mark
ASCI Applications

Mr. Donald Dudley
Mr. Greg Astfalk
Mr. Russ Adamchak
Mr. Ron Matlock

Mr. Richard Kaufmann
Mr. Jamshed Mirza

Ms. Debra Goldfarb
Dr. Earl Joseph II
Mr. James Rottsolk
Rene Copeland
Charles Hayes
Charles Weinhfocker

Mr. Randy Christensen,
Lawrence Livermore
National Laboratories

ANNEX B. TASKING MEMORANDUM



ACQUISITION AND
TECHNOLOGY

THE UNDER SECRETARY OF DEFENSE

3010 DEFENSE PENTAGON
WASHINGTON, DC 20301-3010

15 NOV 1999

MEMORANDUM FOR CHAIRMAN, DEFENSE SCIENCE BOARD

SUBJECT: DoD Super Computing Needs

Recent commercial developments in the super computing industry have highlighted DoD needs in this specialized community. It is therefore both timely and important for the Defense Science Board (DSB) to place a special focus on this critical technology.

The rapidly changing super computing technology offers DoD an opportunity to investigate new alternatives to existing capability. Thus, we would like the DSB effort to focus on alternative super computing technologies especially in the areas of distributed networks and multi-processor machines. The TF should pay particular attention to affordability of new technologies and associated risks.

Towards that end, please ensure that the Chairman of the DSB Task Force on Defense Software establishes an appropriate sub-group to address DoD super computing needs, especially as related to the field of cryptography requirements.

The Task Force shall have access to classified information needed to develop its assessment and recommendations.

Further request that the sub-group's findings and conclusions be provided to me in the form of a letter report at the earliest possible opportunity.


Jacques S. Gansler

